Save Energy Now in High Performance Computer Centers



Super Computing 08 November 20, 2008

Dale Sartor

Lawrence Berkeley National Laboratory

DASartor@lbl.gov







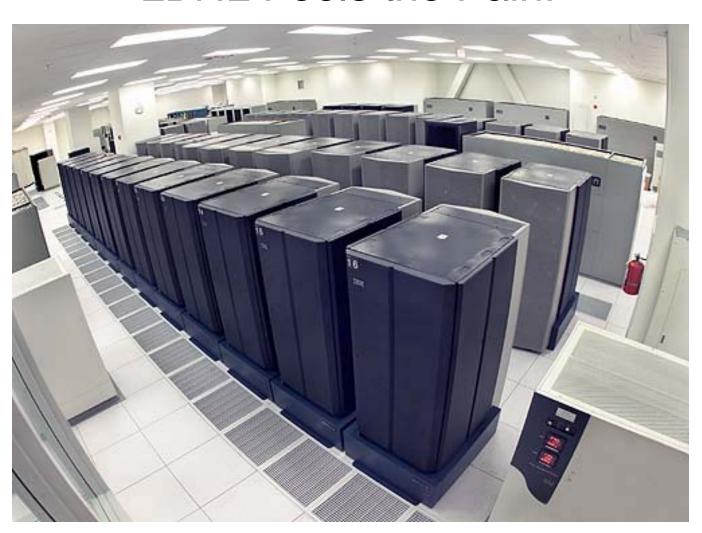
Agenda:

- Trends
- Benchmarking
- Best Practices
- LBNL's experience

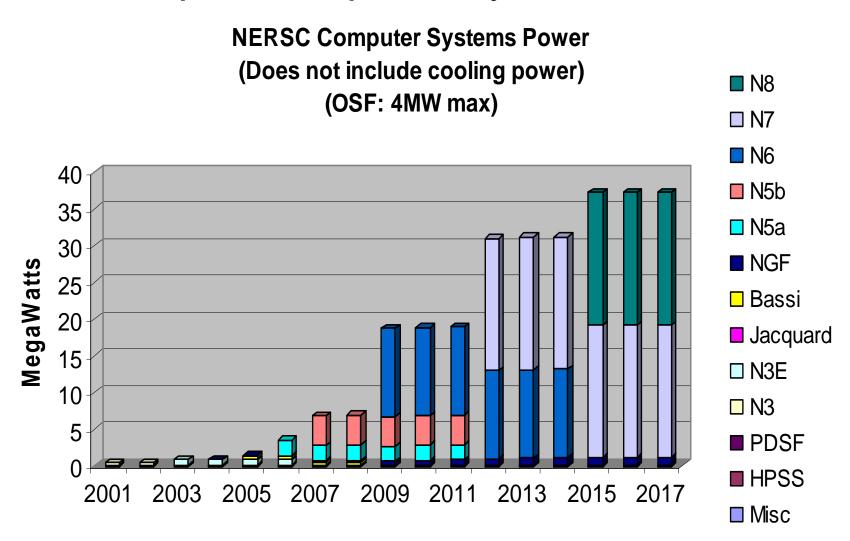
HPC and Super Computer Power Trends:

- Super computer power consumption is expected to double every 3-4 years, while functionality is expected to increase four times every three years
 - Functionality is increasing at least twice as fast as power, but power is a growing concern
- Super computer power growth will be much faster than general data center power growth (20+% vs. 12% as estimated in the EPA report to Congress)
- HPC server market sales have increased 19% annually in the last four years (IDC 2008)
- Supercomputers are 26% of the HPC market and 5% of the overall server market (IDC 2008)
- HPC sales to government labs was \$1.6B and Defense was another \$1B; combined they represent 25% of the HPC market (IDC 2008)

LBNL Feels the Pain!



LBNL Super Computer Systems Power:



Lifetime electrical cost will soon exceed cost of IT equipment in typical data centers, but IT equipment load can be controlled (not all methods applicable to super computers):

- Consolidation
- Server efficiency
 - Flops per watt
 - Efficient power supplies
- Software efficiency (Virtualization, MAID, etc.)
- Power management
 - Low power modes
- Redundant power supplies
- Reducing IT load has a multiplier effect
 - Equivalent savings +/- in infrastructure

Potential Benefits of Improved Data Center Energy Efficiency:

- 20-40% savings typically possible
- Aggressive strategies can yield better than 50% savings
- Extend life and capacity of existing data center infrastructures
- But is my center good or bad?



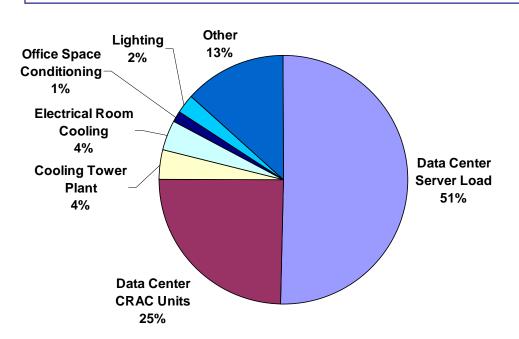
Benchmarking for Energy Performance Improvement:

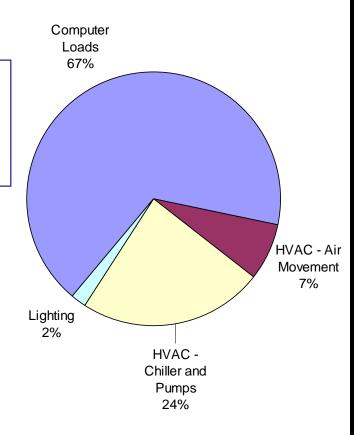
- Energy benchmarking can allow comparison to peers and help identify best practices
- LBNL conducted studies of over 30 data centers:
 - Found wide variation in performance
 - Identified best practices



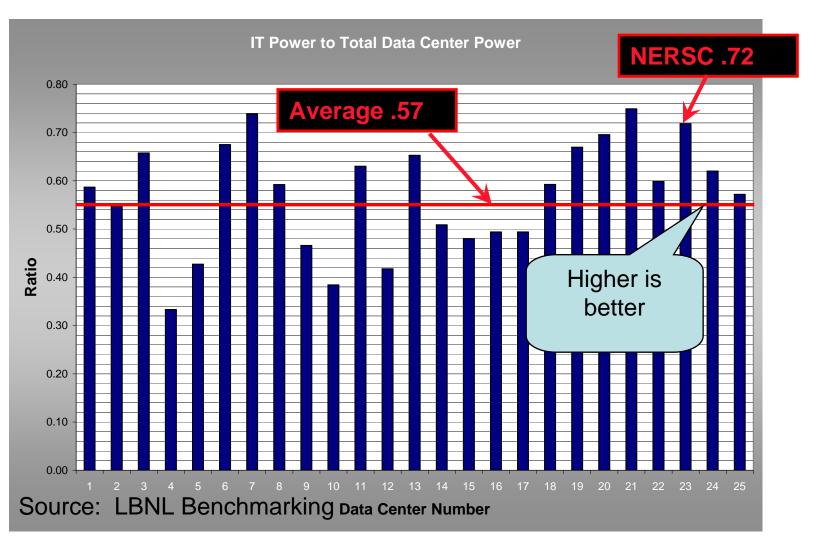
Your Mileage Will Vary

The relative percentages of the energy actually doing computing varied considerably.

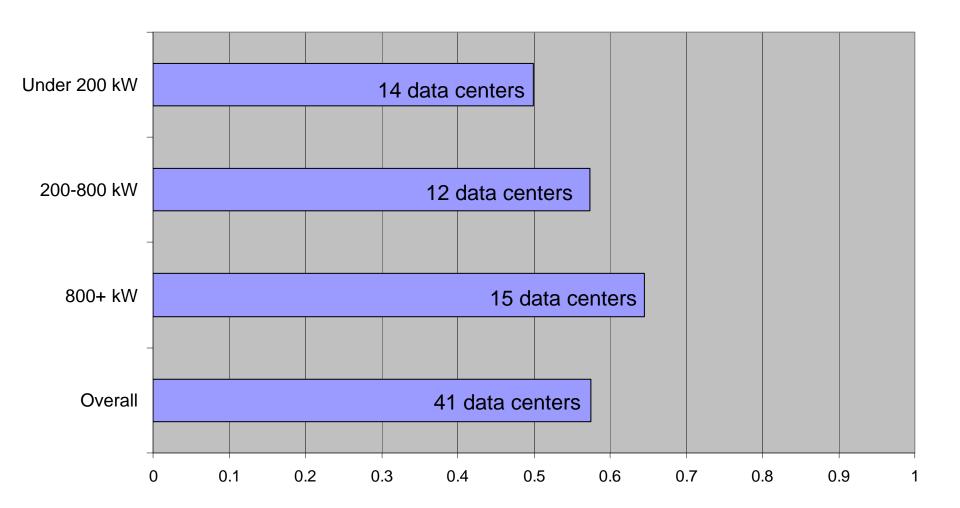




High Level Metric— Data Center Infrastructure Efficiency (DCiE) Ratio of Electricity Delivered to IT Equipment to Total



Average Data Center Infrastructure Efficiency (DCiE) for DOE Data Centers



11

Save Energy Now On-line profiling tool: "Data Center Pro"



INPUTS

Description

Utility bill data

System information

IT

Cooling

Power

On-site gen



OUTPUTS

Overall picture of energy use and efficiency

End-use breakout

Potential areas for energy efficiency improvement

Overall energy use reduction potential

/! with the Quick PEP tool. Use one of the links below to view the online tutorial, start a new case or reg

Quick PEP will provide you with a customized, printable report that shows the details of energy purchases at your plant, how energy is consumed at

your plant, potential cost and energy savings at your plant, and a list of

 If you're new to Quick PEP, view the It's the easiest way to learn how to use

allow to get you started saving energy at your plant.

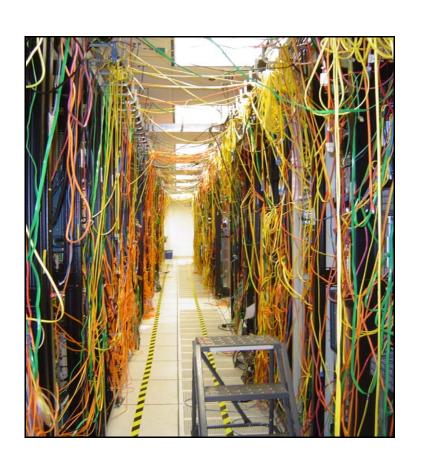
Start a New Case or Open a Case - If you've got all oready, just cick here to start a new Quick PEP case. O saved case.





Using benchmark results to find best practices:

- Air management
- Right-sizing
- Central plant optimization
- Efficient air handling
- Liquid cooling
- Free cooling
- Humidity control
- Improve power chain
- On-site generation
- Design and M&O processes



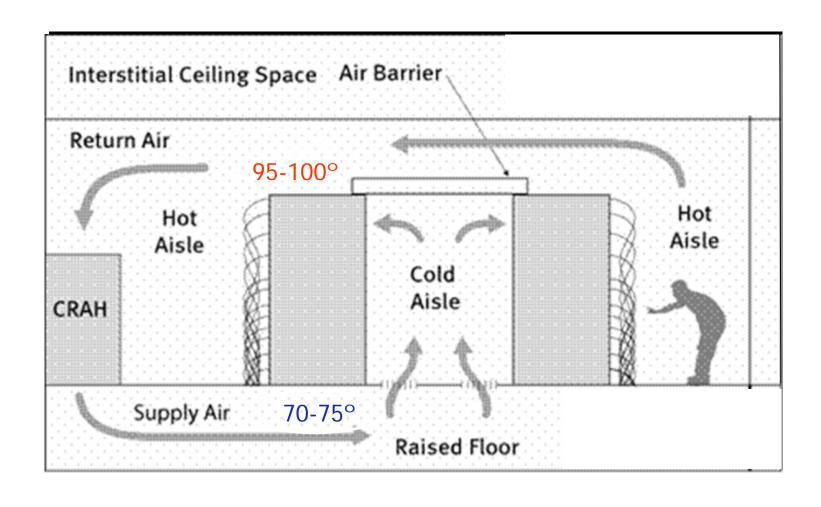
Applying Best Practices at LBNL:

- Partnership between CIO, CS, and energy efficiency researchers
- Existing data centers relatively efficient
 - ➤ NERSC: .72 kW-IT/kW-total (taking advantage of central plant)
 - > 50B-1275: .59 .64 kW-IT/kW-total (tower cooled CRACs)
- Increased efficiency frees up needed "capacity"
- Leveraging data centers as test beds to create an impact beyond Berkeley Lab
- Working with vendors to develop new products and strategies

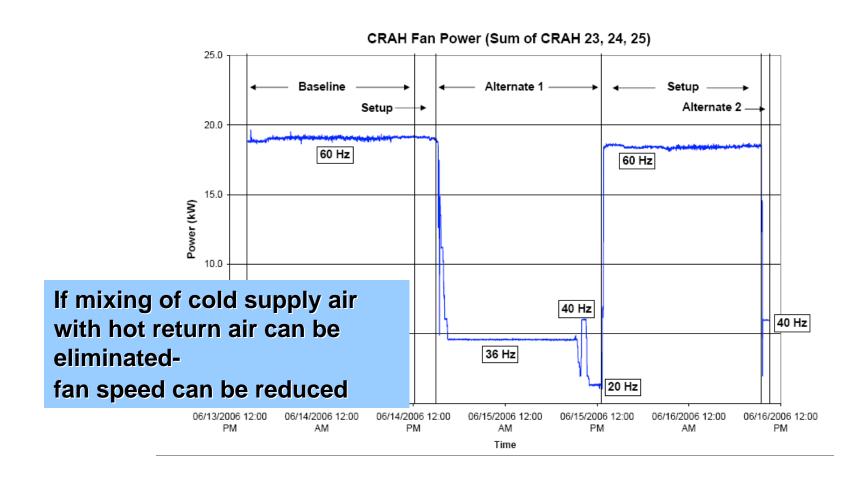
Air Management:

- Typically, much more air is circulated through computer room air conditioners than is required
- Air mixing and short circuiting leads to:
 - Low supply temperature
 - Low Delta T
- Improve isolation of hot and cold "aisles"
 - Reduce fan energy
 - Improve air-conditioning efficiency
 - Increase cooling capacity

Best Scenario— Isolate Cold and Hot

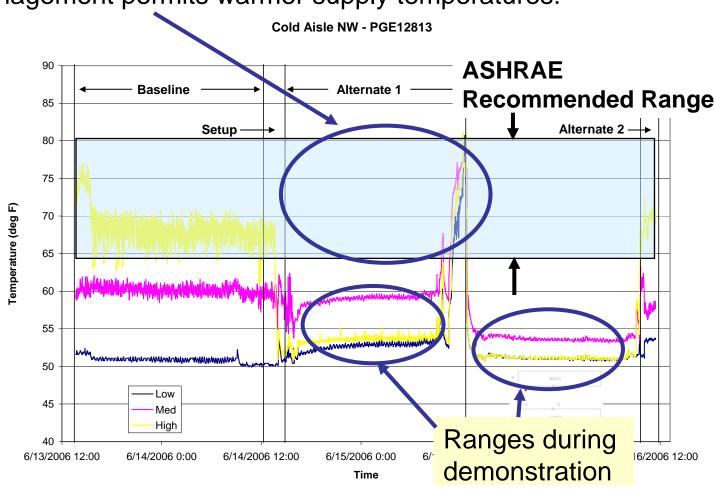


NERSC Test: Temporary Isolation Fan Energy Savings – 75%

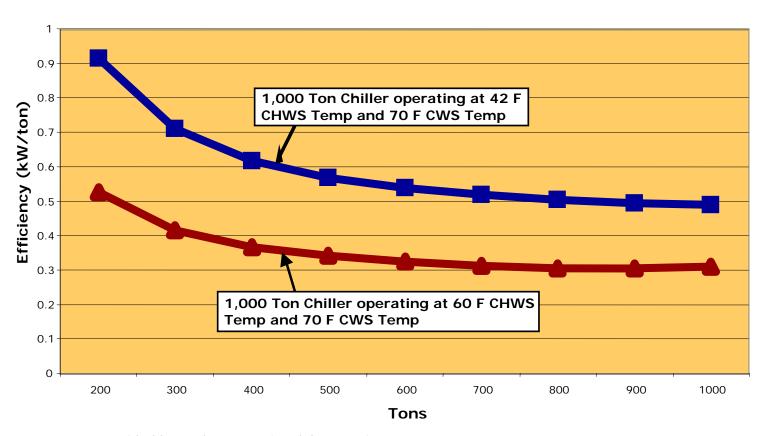


LBNL NERSC air management demonstration

Better airflow management permits warmer supply temperatures!

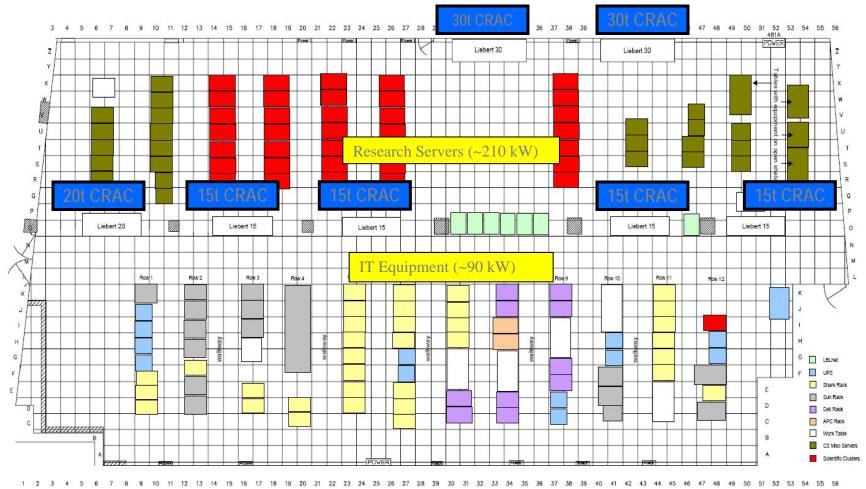


Cooling System Efficiencies Under Different Conditions:



Data provided by York International Corporation.

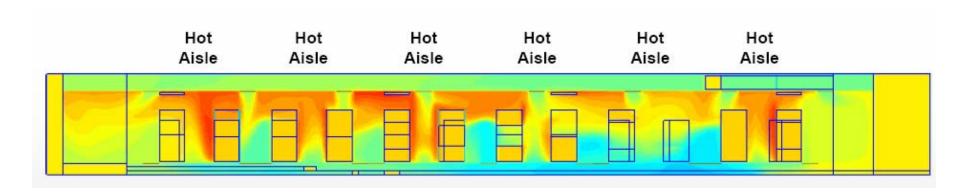
Air Management Improvements in LBNL 50B-1275 Data Center (layout September 2007)

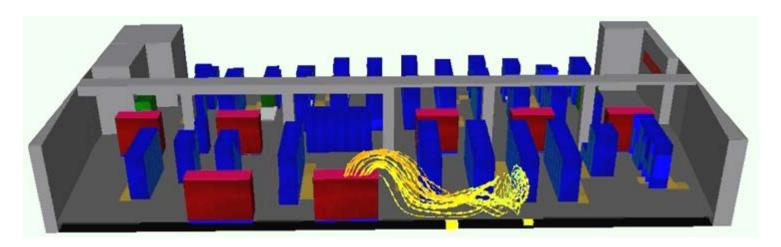


Air Management Assessment Effort:

- Performed CFD
- Deployed wireless monitoring system
- Identified opportunities for improvement
 - Enforce hot aisle/cold aisle arrangement
 - Use blanking panels
 - Improve airflow and under floor pressure by tuning floor tiles (e.g. reduced number of perforated tiles)
 - Reduce mixing and short circuits
 - convert overhead plenum from supply to hot-air return
 - CRAC intakes extended into overhead

Visualization Using CFD:

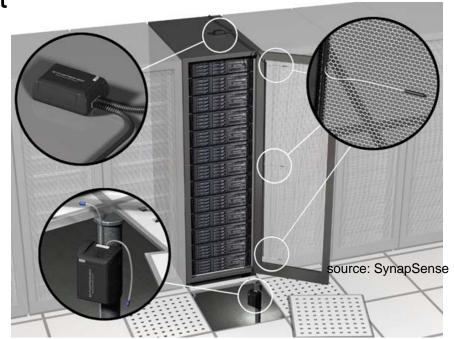




Images: ANCIS

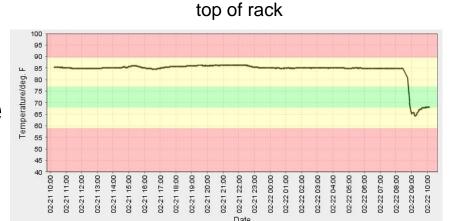
Wireless Environmental Monitoring:

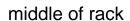
- Most operators lack visibility into data center environment
 - Can't effectively manage the facility
- SynapSense wireless sensor network installed
 - 300 monitoring points (temperature, humidity, under-floor pressure, current)
- Air management now based on empirical data, not intuition

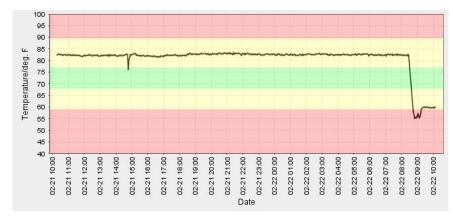


Results: Blanking Panels

- One blanking panel added and temperature dropped ~20°
- impact of other best practices confirmed
 - eliminate leaks in floor
 - improve air management





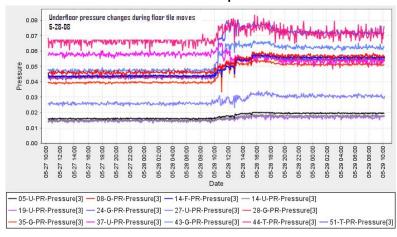


Results: Tune Floor Tiles

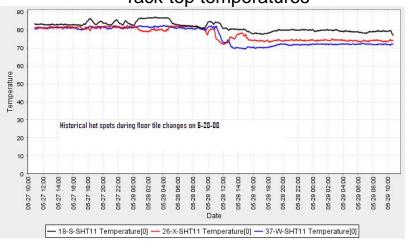


- Too many permeable floor tiles
- if airflow is optimized,
 - under-floor pressure
 - rack-top temperatures
 - data center capacity grows
- Measurement and visualization assisted tuning process

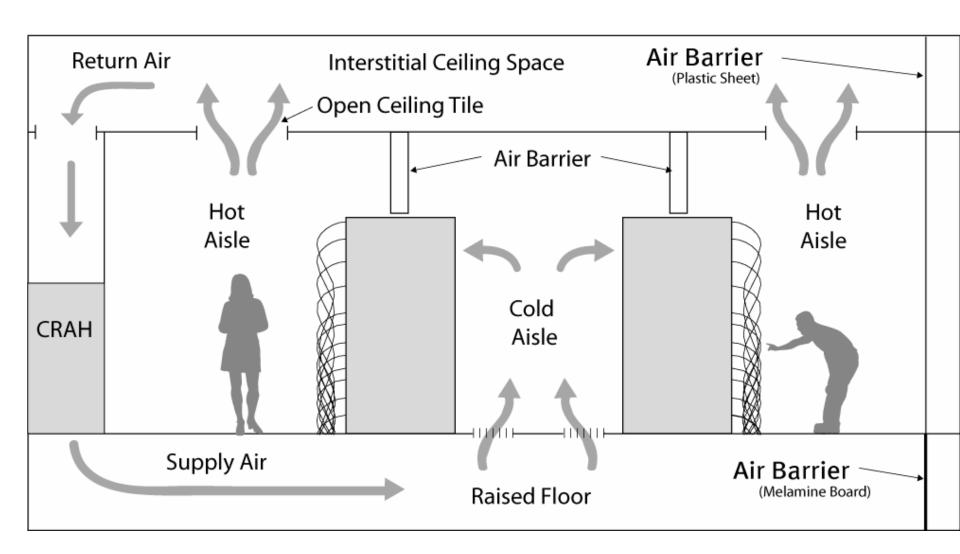
under-floor pressures



rack-top temperatures



Improve Air Management:



Improve Air Management:

- Overhead plenum converted from supply to hot-air return
- CRAC intakes extended into overhead
- Return registers placed over hot aisle



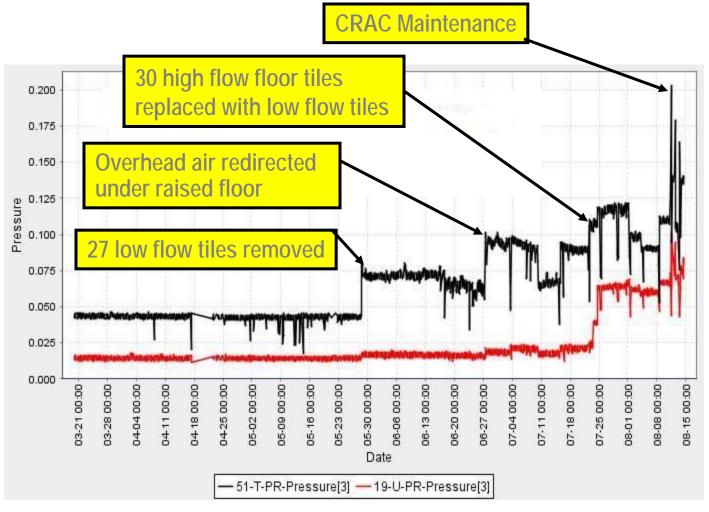
Before



After



Documenting Progress



5 months of floor pressure management

Results: Air Management Project (so far)

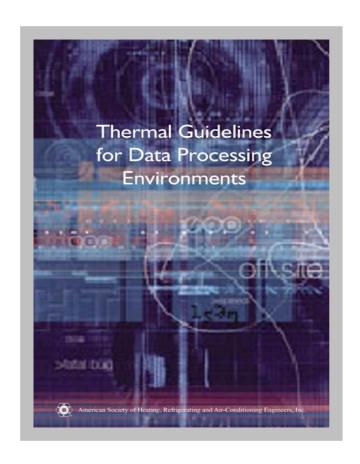
- 21% increase (~100kW) in IT load
- CRAC unit set points 3°F warmer
- Fewer hot spots
- (1) 15 ton CRAC unit turned off
- (1) extra 15 ton CRAC unit on-line but redundant
- Wireless sensor network enables operations to monitor and fine tune changes
- Next step:
 - Curtains to improve isolation
 - Integrated CRAC unit control upgrade

Best air management practices:

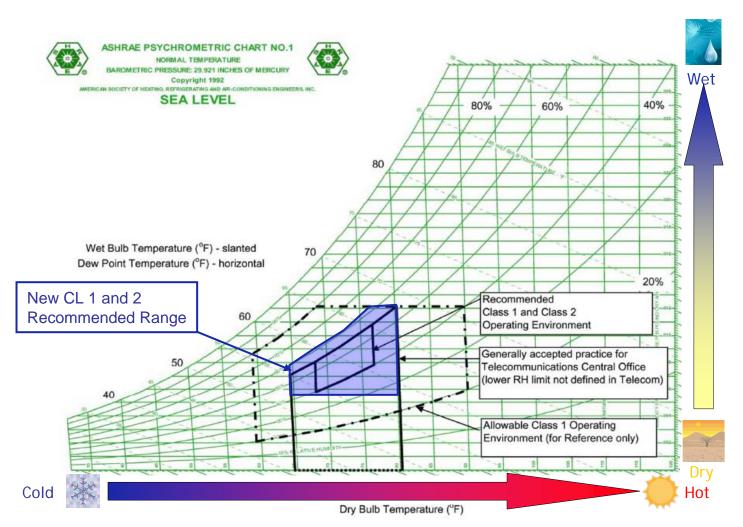
- Arrange racks in hot aisle/cold aisle
- Plug floor leaks and provide blanking plates in racks
- Match server airflow by aisle
- Get variable speed or two speed server fans
- Provide variable airflow fans for AC units
- Draw return from as high as possible or utilize return plenum
- Provide isolation of hot and cold spaces
- Consider overhead supply
- Use CFD and monitoring to inform design and operation

Environmental Conditions

Use ASHRAE
 Recommended and
 Allowable ranges of
 temperature and
 humidity



Design conditions at the inlet to IT equipment



Humidity Ratio Pounds Moisture per Pound of Dry Air

Use Free Cooling:

- Water-side Economizers
 - No contamination question
 - Can be in series with chiller
- Outside-Air Economizers
 - Can be very effective (24/7 load)
 - Must consider humidity
- Let's get rid of chillers in data centers

Free Cooling - Liquid Based

- Infrastructure installed in 2008 for liquid cooling in 50B-1275
- Cooled with tower only or chiller assisted
 - Both options
 significantly better
 than existing liquid
 cooled (DX) CRAC
 units







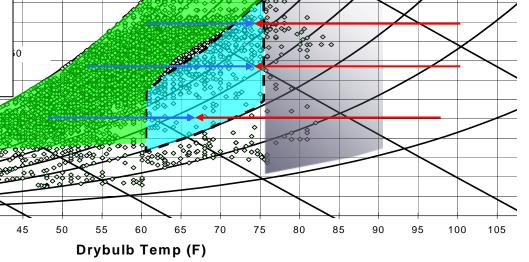
Free Cooling - Outside Air Based

Annual Psychrometric Chart of Oakland, CA
(relative humidity lines are stepped by 10%,



- 2. Green can become blue mixing return and outdoor air
- Most of the conditions below and right of blue can be satisfied w/ evaporative cooling
- 4. Hot and humid hours will enter the "allowable" range or require compressor air conditioning

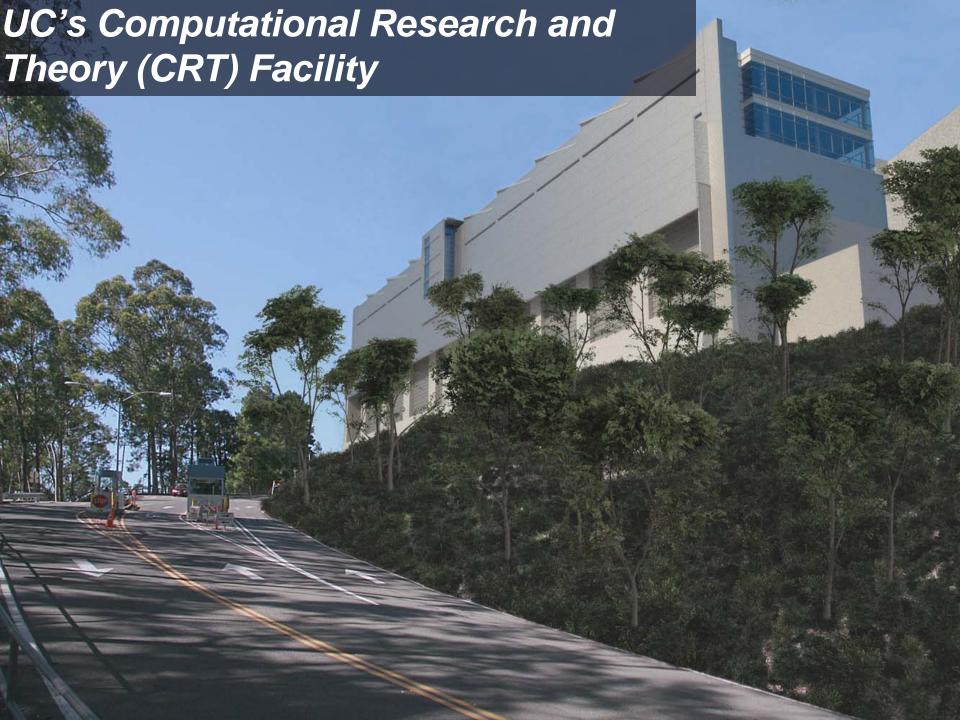
WB = 30



wetbulb lines by 10 degrees F)

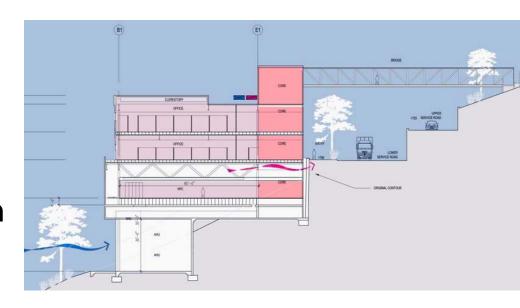
WB = 70

WB = 60



System Design Approach:

- Air-Side Economizer (93% of hours)
- Direct Evaporative Cooling for Humidification/ precooling
- Low Pressure-Drop Design (1.5" total static)



Hours of Operation

Mode 1	
Mode 2	
Mode 3	
Mode 4	
Mode 5	
total	

100% Economiser	2207	hrs
OA + RA	5957	hrs
Humidification	45	hrs
Humid + CH cooling	38	hrs
CH only	513	hrs
	8760	hrs

Water Cooling: Four-pipe System

- Closed-loop treated cooling water from cooling towers (via heat exchanger)
- Chilled water from chillers
- Allows multiple temperature feeds at server locations
- Headers, valves and caps for modularity and flexibility

Predicted CRT Performance:

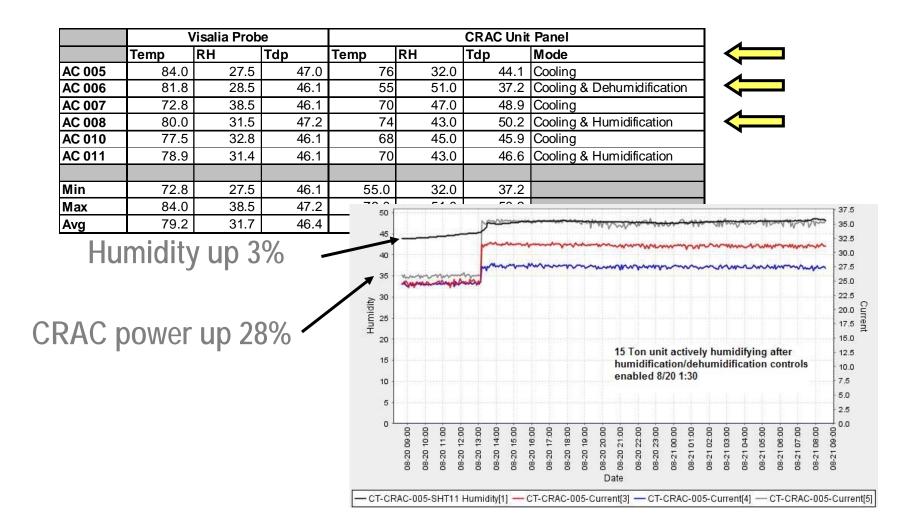
- DCIE of 0.95 based on annual energy
- DCIE of 0.88 based on peak power



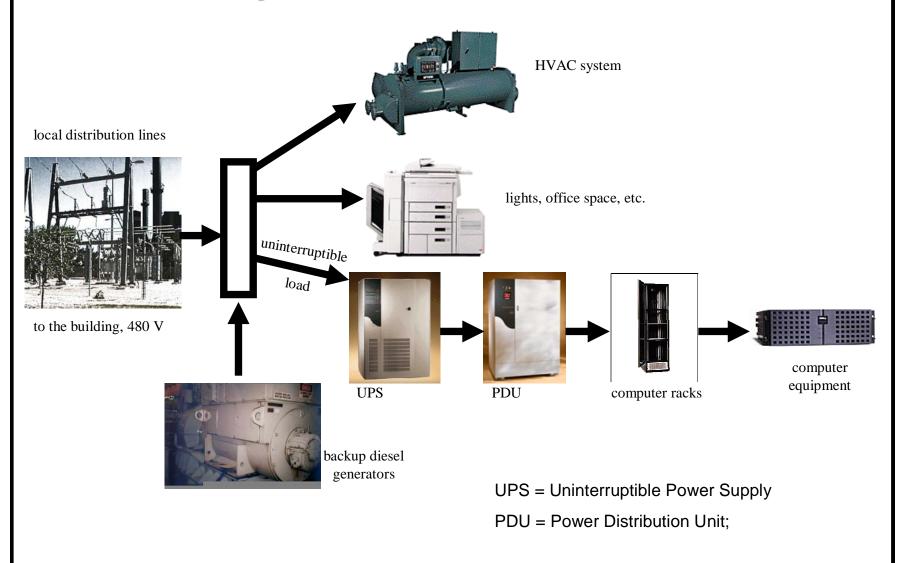
Improve Humidity Control:

- Eliminate inadvertent dehumidification
 - Computer load is sensible only
 - Medium-temperature chilled water
 - Humidity control at make-up air handler only
- Use ASHRAE allowable RH and temperature
- Eliminate equipment fighting
 - Coordinate controls on distributed AHUs

The Cost of Unnecessary Humidification in 50B-1275



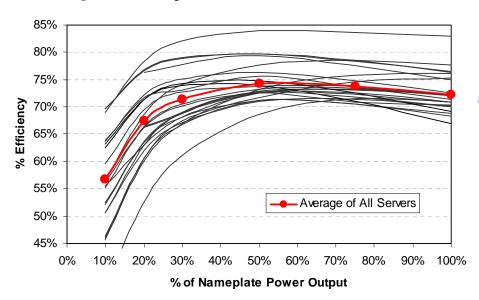
Electricity Flows in Data Centers



Improving the Power Chain:

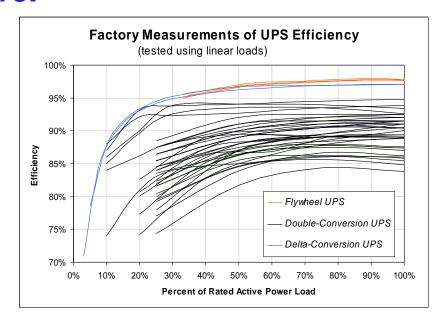
- Increase distribution voltage
 - NERSC going to 480 volts to the racks
- Improve equipment power supplies
- Improve UPS
 - LBNL uses minimal UPS
 - Selection of delta-conversion or doubleconversion with bypass to minimize losses
- DC distribution

Specify Efficient Power Supplies and UPSs



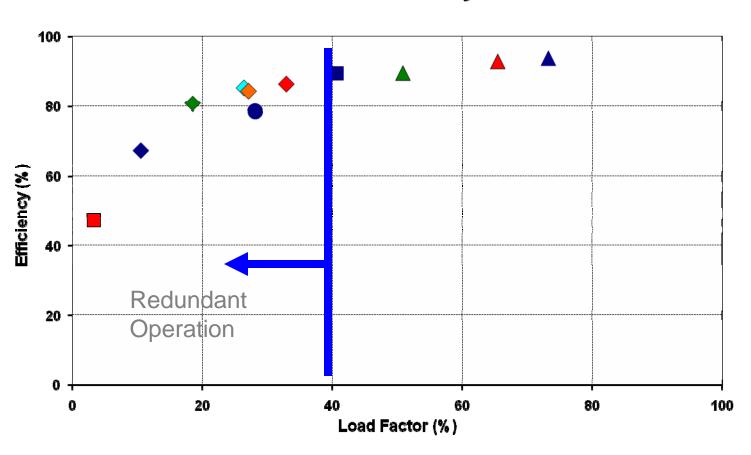
Power supplies in IT equipment generate much of the heat. Highly efficient supplies can reduce IT equipment load by 15% or more.





Measured UPS Efficiency

UPS Efficiency



Redundancy

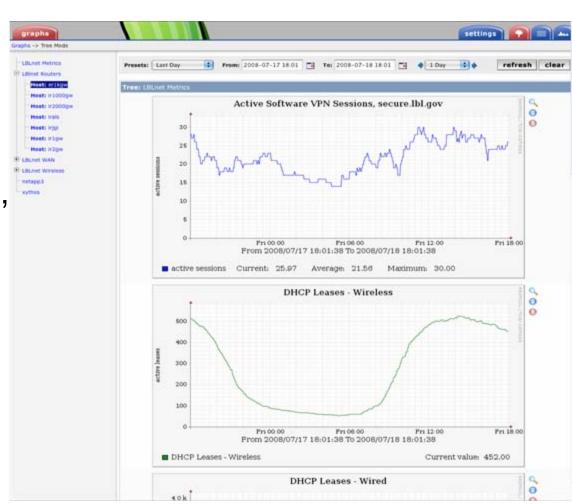
- Understand what redundancy costs is it worth it?
- Different strategies have different energy penalties (e.g. 2N vs. N+1)
- Redundancy in electrical distribution puts you down the efficiency curve
- LBNL minimizes use of redundant power supplies and size of UPS

Improve Design and Operations Processes:

- Get IT and Facilities people to work together
- Use life-cycle total cost of ownership analysis taking all costs into account (e.g. infrastructure)
- Document design intent and provide training
- Benchmark and track existing facilities
- Re-commission regularly as part of maintenance
- Continuously monitor performance

The Importance of Visualization

- Systems & network administrators have tools for visualization
- Useful for debugging, benchmarking, capacity planning, forensics
- Data center facility managers have comparatively poor visualization tools

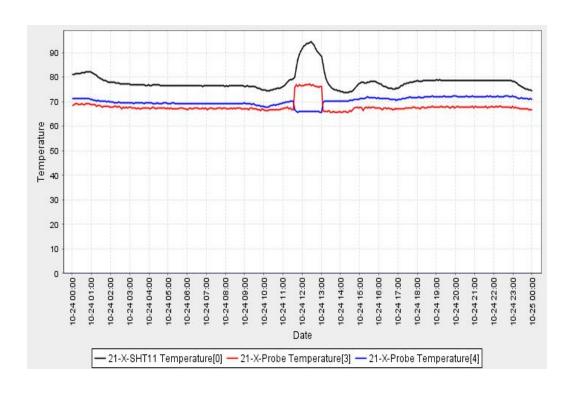


LBNL Wireless Monitoring System:

- 300 point SynapSense wireless monitoring
 - Temperature, humidity, under-floor pressure, CRAC current
 - 1-2 day installation, easy redeployment
- For the first time, we have a detailed understanding of environmental conditions
 - Real-time and historical data
 - Remote console and alert notification
 - Quick reports and graphs from underlying database
 - Modules can also measure liquid flow, liquid presence, and particle count

Learning from the sensors

Rack top, middle and bottom temperature when cold aisle air flow obstructed





Next Steps

- Integration of monitoring system with controls (Automatic staging of CRACs and Demand based resets of pressure and temperature)
- 'Live Imaging' heat-map animations
- Real time DCiE (data center infrastructure efficiency) calculation

Advice:

- Put together team of IT, Facilities and experienced consultants
- Benchmark energy and environmental performance
- Instrument and convert data to information
- Eat your spinach (blanking panels, leaks, CRAC maintenance)
- Keep an eye on emerging technologies (flywheel UPS, rack-level cooling, DC power)

Top best practices identified through benchmarking

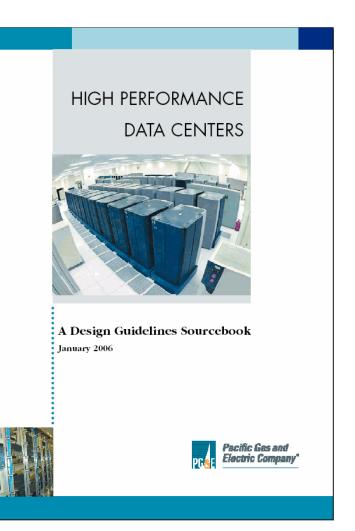
HVAC - Air Delivery - Water Systems		Facility Electrical Systems	IT Equipment	Cross-cutting / misc. issues
Air management	Cooling plant optimization	UPS systems	Power Supply efficiency	Motor efficiency
Air economizers	Free cooling	Self generation	Sleep/standby loads	Right sizing
Humidification controls alternatives	Variable speed pumping	AC-DC Distribution	IT equip fans	Variable speed drives
Centralized air handlers	Variable speed Chillers	Standby generation		Lighting
Direct liquid cooling				Maintenance
Low pressure drop air distribution				Commissioning/continuous benchmarking
Fan efficiency				Heat recovery
				Redundancies Method of charging for
				space and power Building envelope

Completed

Design Guidelines Are Available

- Design Guides were developed based upon the observed best practices
- Guides are available through PG&E and LBNL websites
- Self benchmarking protocol also available

http://hightech.lbl.gov/datacenters.html



Industrial Technologies Program

- Tool suite & metrics
- Energy baselining
- Training
- Qualified specialists
- Case studies
- Certification of continual improvement
- Recognition of high energy savers
- Best practice information
- Best-in-Class guidelines

EPA

- Metrics
- Server performance rating & ENERGY STAR label
- Data center performance benchmarking

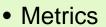


Federal Energy Management Program

- Best practices showcased at Federal data centers
- Pilot adoption of Best-in-Class guidelines at Federal data centers
- Adoption of to-be-developed industry standard for Best-in-Class at newly constructed Federal data centers

Industry

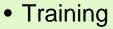












- Best practice information
- Best-in-Class guidelines
- IT work productivity standard











Links to Get Started

DOE Website: Sign up to stay up to date on new developments www.eere.energy.gov/datacenters

Lawrence Berkeley National Laboratory (LBNL) http://hightech.lbl.gov/datacenters.html

LBNL Best Practices Guidelines (cooling, power, IT systems) http://hightech.lbl.gov/datacenters-bpg.html

ASHRAE Data Center technical guidebooks http://tc99.ashraetcs.org/

The Green Grid Association – White papers on metrics http://www.thegreengrid.org/gg_content/

Energy Star® Program

http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency

Uptime Institute white papers www.uptimeinstitute.org

Join network to share information and Pull market towards higher efficiency products:

Contact Information:

Dale Sartor, P.E.

Lawrence Berkeley National Laboratory

Applications Team MS 90-3111 University of California Berkeley, CA 94720

DASartor@LBL.gov

(510) 486-5988

http://Ateam.LBL.gov

